THIN-FILM INTRACORTICAL RECORDING MICROELECTRODES

Quarterly Report #4

(Contract NIH-NINDS-NO1-NS-0-2329)

July - September 2001

Submitted to the

Neural Prosthesis Program
National Institute of Neurological Disorders and Stroke
National Institutes of Health

by the

Center for Wireless Integrated MicroSystems
Department of Electrical Engineering and Computer Science
The University of Michigan
Ann Arbor, Michigan
48109-2122

October 2001

Thin-Film Intracortical Recording Microelectrodes

Summary

During the past quarter, we have proceeded to develop chronically-implantable multi-channel recording probes in a number of areas. Work has begun to decrease the shank widths on the probes to much smaller dimensions (<20µm) while accommodating recording sites in the 600µm to 800µm range. A process based on polysilicon refill of a deep trench along the sides of the probe is being explored.

We have just completed the lastest fabrication run of active probes. Circuit thresholds are on-target at +0.9V (nMOS) and –0.6V (pMOS). Probes containing 3µm-thick gold shields over the circuitry have also been fabricated in preparation for chronic implant studies of the active structures. Some of the active probes contain silicon ribbon cables to allow them to be implanted without a surface platform, and for the first time such ribbons have been released entirely from the front using trenches formed along the sides of the cable. Also during the past quarter, the inputs from active probes have been stabilized with respect to their dc potentials using subthreshold shunt MOS transistor clamps. Without the clamps, the dc input level wanders ±1mV over short times and ±10-40mV over longer periods. With the clamps activated, the dc input level is stable to better than ±50µV, while the lower ac cutoff frequency varies from about 30Hz to less than 3Hz, depending on bias. The use of a variable bias on the clamping devices would provide a means for varying the bandpass of the recording system depending on application.

With the dc baseline of the probes stable over time, it should be possible to implement systems in which the recorded spikes are thresholded and either the simple occurrence of a spike (together with its site address) are passed off the probe or a digital word giving its amplitude and address are transmitted. This can be expected to save large amounts of bandwidth compared to full analog signal transmission and sets the stage for prosthetic systems in which signal interpretation is done in-vivo. A platform chip performing spike thresholding has been designed and fully simulated to demonstrate its functional operation. The chip will be fabricated through MOSIS and used on platform-mounted 3D arrays of PIA-3 64-site 8-channel probes.

Finally, we have obtained the latest version of the telemetry interface back from fabrication and have begun testing it. The chip features a new front-end interface block together with a sigma-delta demodulator and a transmitter. The on-chip regulators consume 167µW for in input voltage of 12V and offer regulation of 1mV/V with a ripple rejection ratio of >30dB. The sigma-delta modulator consumes 230µW and produces an output bit rate at 2MHz. The total telemetry interface chip dissipates less than 900µW from 5V and measures 2.2mm x 2.2mm. We are moving to complete testing on both the PIA-2 chips and the telemetry interface and hope soon to test the probes in a wireless configuration in-vivo.

Thin-Film Intracortical Recording Microelectrodes

1. Introduction

The goal of this program is the realization of batch-fabricated recording electrode arrays capable of accurately sampling single-unit neural activity throughout of volume of cortical tissue on a chronic basis. Such arrays will constitute an important advance in instrumentation for the study of information processing in neural structures and should be valuable for a number of next-generation closed-loop neural prostheses, where stimuli must be conditioned on the response of the physiological system.

The approach taken in this research involves the use of solid-state process technology to realize probes in which a precisely-etched silicon substrate supports an array of thin-film conductors insulated above and below by deposited dielectrics. Openings in the dielectrics, produced using photolithography, form recording sites which permit recording from single neurons on a highly-selective basis. The fabrication processes for both passive and active (containing signal-processing circuitry) probe structures have been reported in the past along with scaling limits and the results of numerous acute experiments using passive probes in animals. In moving to chronic implant applications, the major problems are associated with the preserving the viability of the sites in-vivo (preventing tissue encapsulation of the sites) and with the probe output leads, both in terms of their number and their insulation. The probe must float in the tissue with minimal tethering forces, limiting the number of leads to a few at most. The encapsulation of these leads must offer adequate protection for the megohm impedance levels of the sites while maintaining mechanical lead flexibility.

Our solution to the lead problem has involved two steps. The first has been to embed circuitry in the probe substrate to amplify and buffer the signals and to multiplex them onto a common output line. Using this approach, signal levels are increased by factors of over 100, impedance levels are reduced by three to four orders of magnitude, and the probe requires only a few leads for operation, independent of the number of recording sites. A high-yield merged process permitting the integration of CMOS circuitry on the probe has been developed, and this circuitry has been designed and characterized. The second step has involved the development of silicon-based ribbon cables, realized using the same probe technology, to conduct the neural signals to the outside world. These cables have shown significant advantages over discrete leads, both in terms of the ease with which chronic implants can be assembled and in terms of the ability of the cables to survive long-term biased soaks in saline. The cables can be built directly into the probes so that they come off of the wafer as a single unit, requiring no joining or bonding operations between them. The cables are also significantly more flexible than previously-used discrete wire interconnects.

This contract calls for the development of active probes for neural recording. A 64-site 8-channel probe with site selection and signal buffering but no multiplexing has been developed (PIA-2B) along with a high-end multiplexed probe that includes gain

(PIA-2). These probes are now being refined and applied to in-vivo applications. Investigations are on-going to better understand site encapsulation, which limits the lifetime of chronic recording structures, and telemetry is being developed to allow the probes to be operated over a wireless link, eliminating the percutaneous connector.

During the past quarter, we have begun work on reducing the size of the recording probe shanks and have completed a fabrication run on a new set of active probes (PIA-2B and PIA-2/-3). We have also continued the development of a wireless interface for the probes. Work in these areas is discussed in the sections below.

4. Scaled Recording Probe Development

Recording studies call for an increasing number of sites, and to minimize the tissue trauma associated with array insertion and residence, it is important that the shanks be as small as possible. In the past, we have demonstrated that shanks can be made much narrower than our standard dimensions (50-80μm), and we are now developing techniques for reducing shank widths considerably. Such "nanoprobes" should minimize tissue damage, allowing recordings from multiple sites within the field of a single cell as well as the transparent spacing of sites throughout a block of relatively-undisturbed neural tissue.

For our standard probes, the substrate is fabricated by selectively boron doping a silicon wafer through oxide masks and then etching the wafer away to the boron etch-stop in an EDP solution. The thickness and cross-section of the substrate are then determined by the diffusion characteristics of boron into silicon. When the width of the mask opening is large, the diffusion process can be characterized as a constant planar-source diffusion. However, as the width of the diffusion window becomes narrower, the diffusion process should be characterized as a line-source diffusion into an infinite volume. In addition, as the width of the diffusion opening decreases, lateral diffusion of boron underneath the mask becomes more important and eventually limits the minimum achievable shank width for a given substrate thickness.

Figure 1 shows a plot of the measured shank thickness as a function of the mask width for boron diffusion, normalized to the thickness for a large mask width, and Fig. 2 shows the measured shank width as a function of the mask width. We can see that the substrate thickness decreases rapidly as the mask opening is reduced below 10~15μm, or as the mask opening becomes comparable to the etch-stop diffusion depth. Some time ago, we developed a probe with a scaled shank width. A vertical RIE etch was used to slice away the lateral boron extension, leaving a probe width set solely by lithography. After the field dielectrics are opened, RIE removed the field portion of the diffusion, slicing vertically to form the sidewalls of the probe. The width of probe here (Fig. 3) is 5μm, the interconnect is 2μm wide, and the probe thickness is 6-10μm. Using a dry etch to define the probe reduces the shank width while retaining a thick substrate, but the very narrow shank leaves little room for the sites. Some means for allowing sites to overhang the shank is needed.

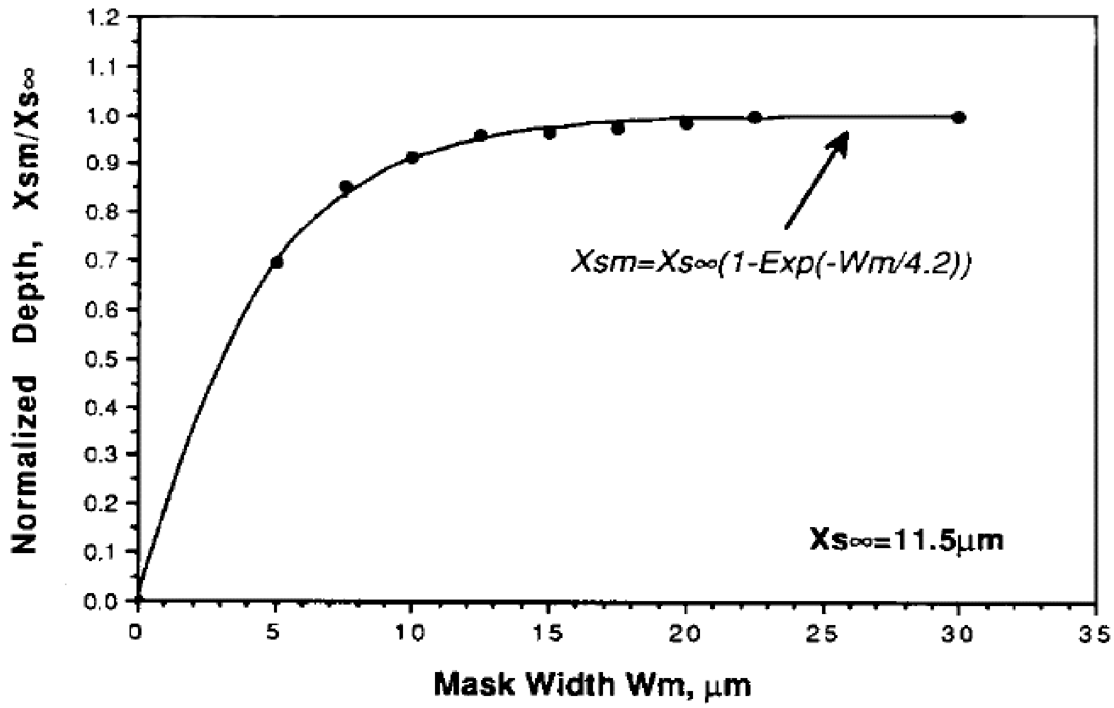

Fig. 1: Plot of the normalized substrate, where x_{sm} is the depth of boron etch-stop.

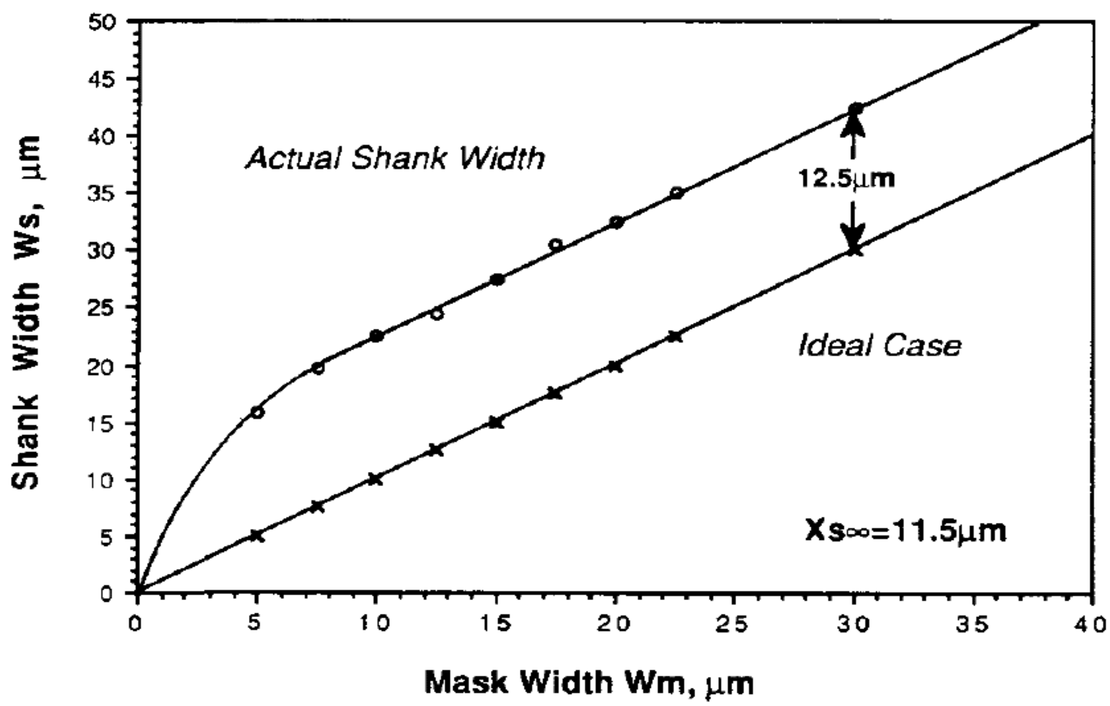

Fig. 2: Plot of the final shank width, Ws, as a function of mask opening, Wm.

From this point of view, we have designed a preliminary "nanoprobe" structure that can minimize the volume displacement of the probe while preserving a reasonable site area. It is noted that our recent studies have indicated advantages in larger sites (600-800μm^2 compared with our more usual 80-100μm^2). The basic process flow of the new design is similar to our usual process. After running the normal process and depositing the top dielectric films, trenches along both sides of the shanks are formed using DRIE (Fig. 4 (a)). The trench is about 2μm wide and 8-10μm deep. Then, polysilicon is deposited over the entire structure, filling the trench (Fig.4 (b)). The polysilicon is dry-etched to remove it from the top surface of the probe and the usual recording (or stimulating) sites are then formed (Fig.4 (c)). Finally, the probe is released. The polysilicon is etched away along with the probe substrate, resulting in the structure shown in Fig.4 (d).

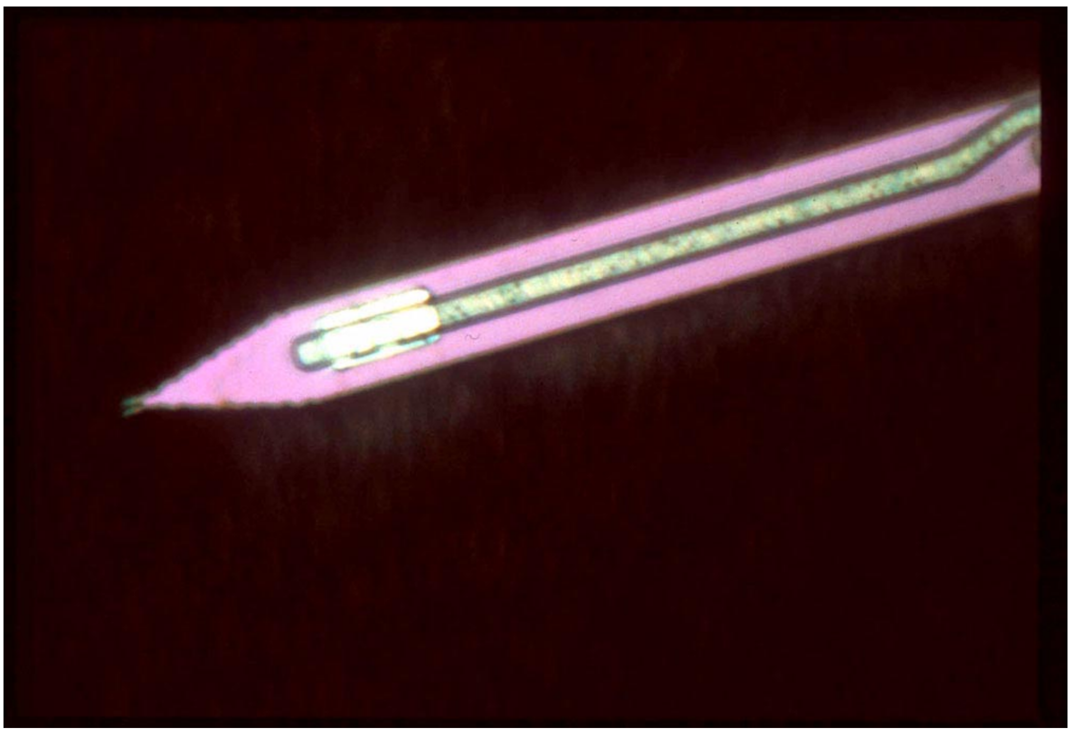

Fig. 3: Scaled recording probe using DRIE

A variety of other probe structures are perhaps possible, including "T"-shaped substrates that allow a wider span for the interconnects as well as the sites, but which minimize the thickness of the substrate with the exception of a "reinforcing bar" down the middle. We are continuing to examine process and structural alternatives and plan to fabricate some of these during the coming term. Strength simulations will also be made. One of the first applications of this technology will be in multi-shank structures where the shanks are perhaps 3μm wide and with which we can record from the field of a single cell in-vivo from multiple portions of the field, both vertically and laterally.

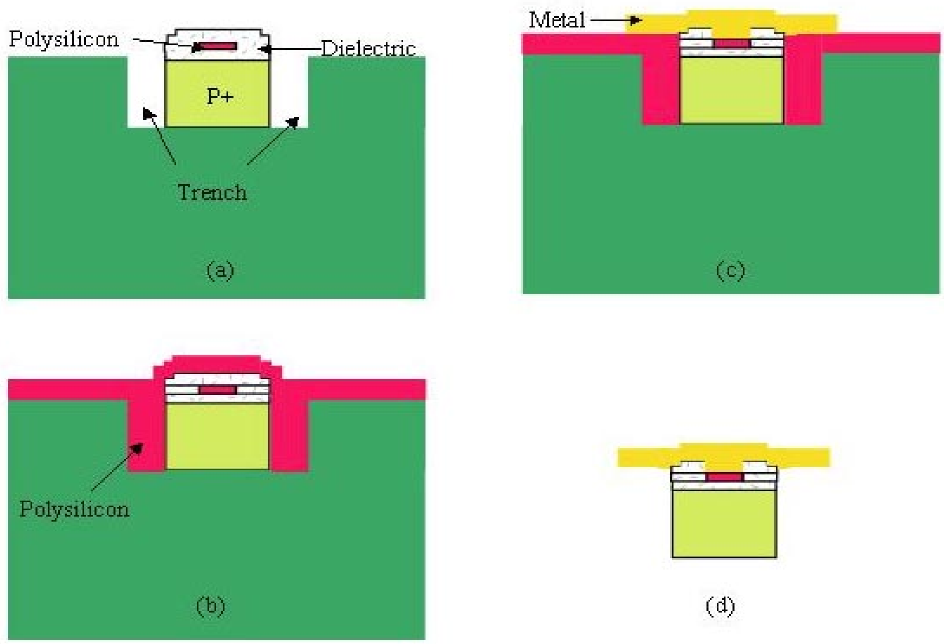

Fig. 4: One possible fabrication procedure for a scaled recording probe.

3. Development of a 64-Site Eight-Channel Non-Multiplexed Recording Probe (PIA-2B)

Fabrication has been completed through circuit metal on the most recent probe designs. Measured process parameters agree very well with expected values, and all circuit blocks are currently being tested. The completed probes will be released from the wafer in the coming weeks. Figures 5 and 6 show the device characteristics.

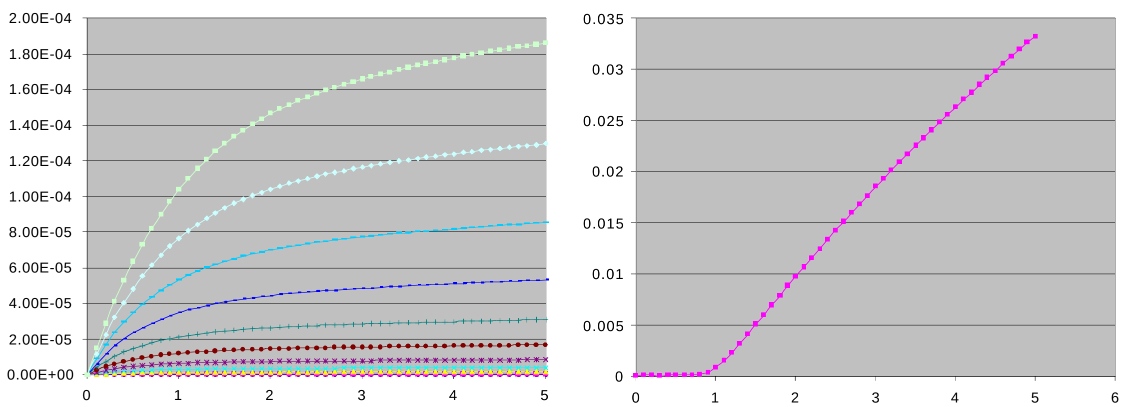

Fig. 5: NMOS transistor I/V curve (left) and square root of drain current versus gate voltage (right.) The x-intercept of the right hand figure gives the NMOS threshold voltage, approximately 0.9V.

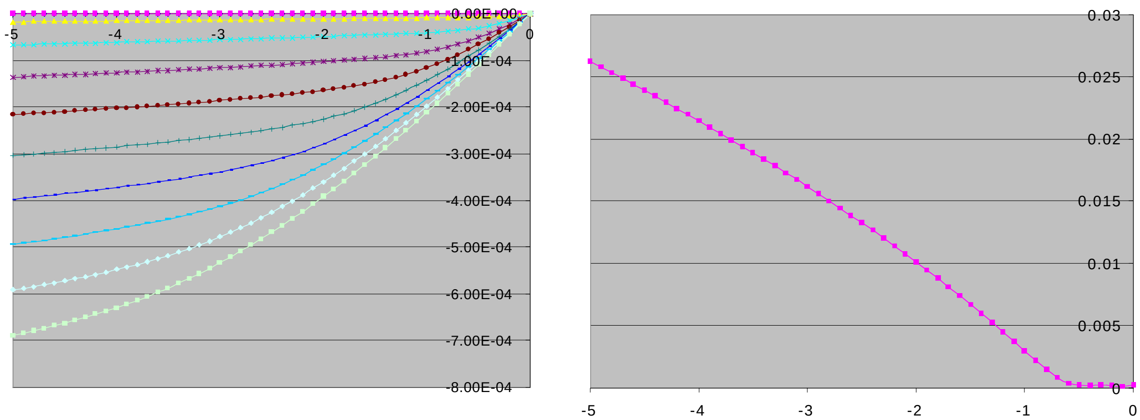

Fig. 6: PMOS transistor I/V curve (left) and square root of drain current versus gate voltage (right.) The x-intercept of the right hand figure gives the PMOS threshold voltage, approximately -0.6V.

Chronic Active Probe Testing

A family of probes has been designed and fabricated to test and improve the ability of Michigan probes to record chronically. These probes will allow the measurement of a mean time to failure for different aspects of probe functionality. Test structures (Fig. 7) have been designed to monitor the etching of circuit metal and polysilicon, to look for theshold shifts under thermal and electrical bias, and to measure leakage currents from metal and polysilicon to saline.

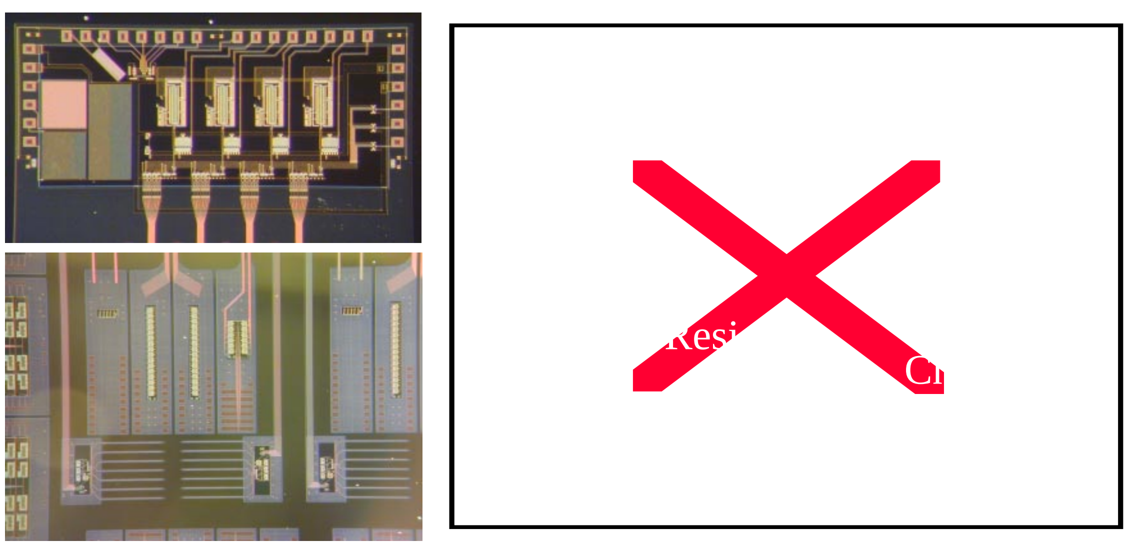

Fig. 7: Probes designed for chronic implantation in guinea pig auditory cortex, left. The probes at bottom left have integrated ribbon cables. Details of the test circuitry, at right.

In addition, the chronic probes have eight recording sites that are front-end-selected onto a single buffered output channel to test overall probe functionality. This probe has been tested exhaustively and is functioning as expected. The on-chip buffer is a unity-gain operational amplifier. This buffer has the advantages of lower output resistance, less DC offset, and gain closer to one than the previously used source follower. The buffer is working as expected, as shown in Fig. 8. In-vitro testing of these probes at body temperature, as well as accelerated testing, will begin in the coming quarter. In-vivo testing will be carried out pending results of the in-vitro tests.

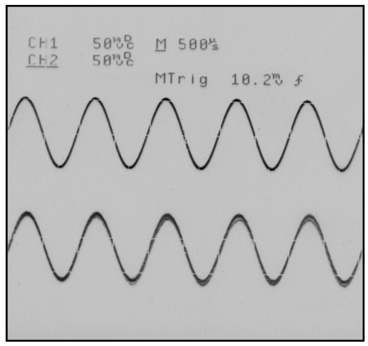

Fig. 8: Oscilloscope trace of the on-chip opamp input/output in the unity-gain configuration.

In order to provide circuit encapsulation for the chronic active probes, a 3µm-thick electroplated gold shield has been added over the circuit area (Fig 9). The gold shield is formed with a single additional mask, and the process is identical to that used to plate gold beam leads for 3D probes. The gold shield has the additional advantage of providing optical shielding for the circuitry.

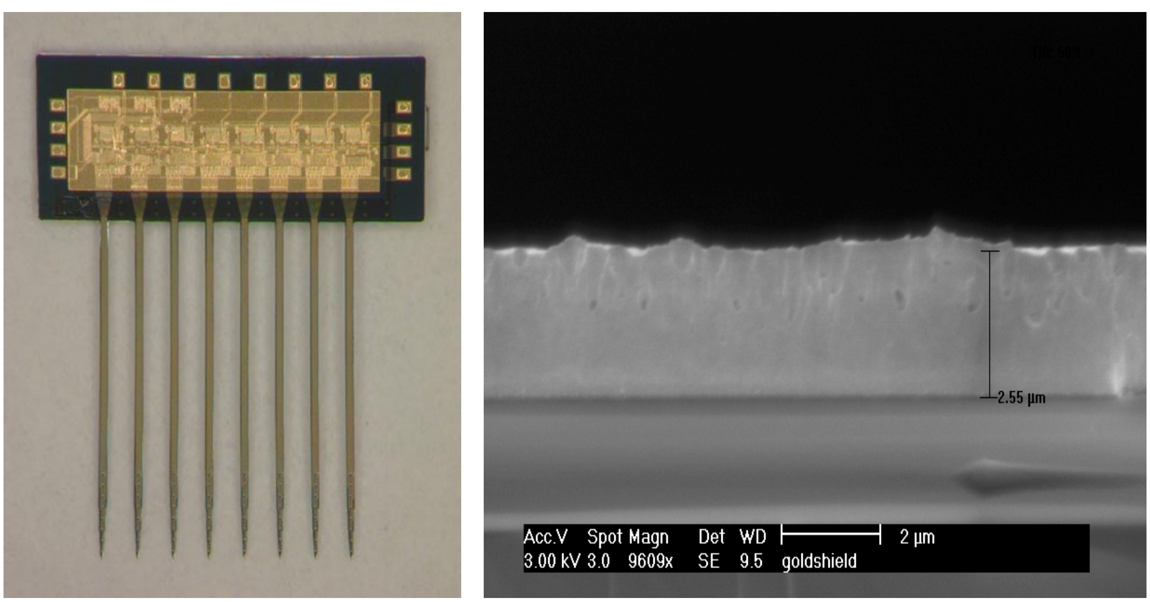

Fig. 9: A probe with an electroplated gold shield over circuit area, left. Cross section of the gold film, right.

Over the past eighteen months, there has been significant progress at Michigan and elsewhere in the ability to obtain chronic in-vivo recordings for periods in excess of a year. The most successful preparation has utilized probes with integrated ribbon cables. For this reason, a version of the chronic active probe was created with a ribbon cable. This creates some difficulty in terms of properly undercutting and etching out the long ribbon while not allowing the circuit area to be attacked. For this reason, a process has been developed for etching deep trenches along the ribbon cable, allowing it to be undercut from the front during EDP release step. This method has produced satisfactory results in test structures and will be used prior to etching out the active probes.

DC Stabilization

During the past quarter, front-end DC stabilization using an integrated shunt sub-threshold bias transistor was demonstrated for the first time. In Fig. 10 below left, an oscilloscope trace of two iridium sites on buffered active probes in saline is given. Uncorrelated variations in the electrode potentials over of several millivolts can be seen at this time scale, which is roughly twenty seconds. At longer time scales, variations can be several tens of millivolts. At right is a trace of the same sites after a shunt transistor is biased into the sub-threshold region. Here, the total peak to peak electrode fluctuation is reduced to the order of 10μV. The resistance of the shunt transistor is dependent on the gate voltage in the subthreshold region, and as a result both the DC level and the lower cutoff frequency of the site are affected by gate bias. Figure 11 gives a family of I/V curves for a shunt transistor as a function of gate bias. The inverse of the slope of each curve is the shunt resistance, and the x intercept is the open circuit potential for the site (i.e., the resulting DC potential of the "clamped" site.) Figure 11 also gives the corner frequency for the high-pass filter resulting from the site capacitance and shunt resistance of the transistor. We can see from this figure that there is a range on the order of 100mV of gate voltage that will provide adequate DC clamping and a sufficiently low cutoff frequency simultaneously. The challenge in integrating this approach into a multi-channel probe will be to design on-chip biasing that will bias multiple channels correctly.

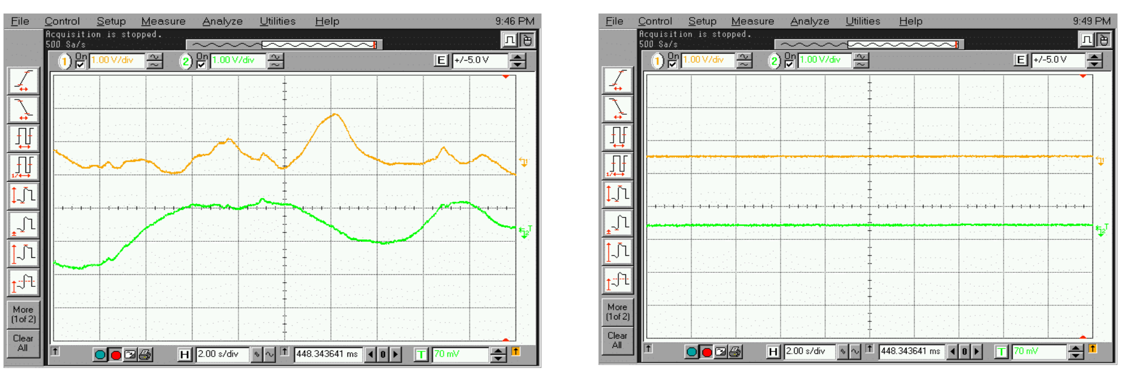

Fig. 10: Baseline variation in electrode potential for two iridium sites in saline (gain is 1000) with and without a sub-threshold biased shunt transistor (left and right, respectively.)

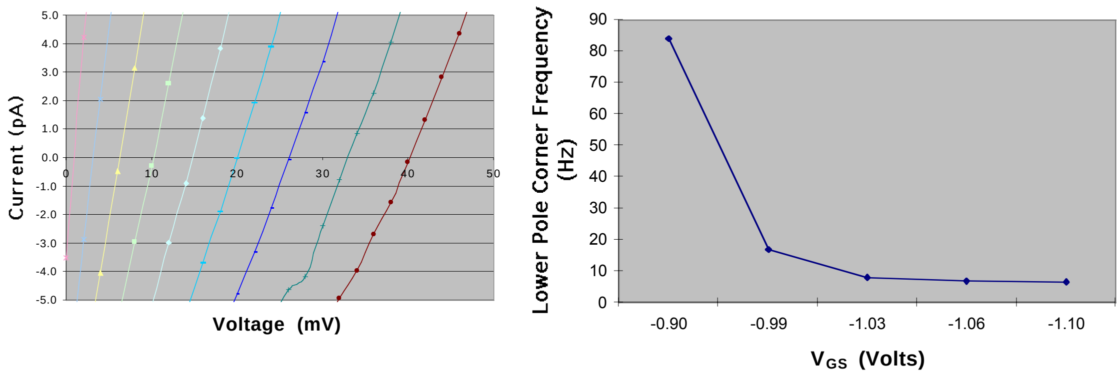

Fig. 11: I-V curve (left) for a shunt transistor at different bias levels, and the resulting low cutoff frequency (right) as a function of gate voltage.

4. Development of On-Platform Signal Processing for Chronic Recording Probes

The development an on-platform ASIC to perform in-vivo spike detection is underway. A behavioral-level verilog program has been written that describes the overall function of the spike-detection unit. A system-level block diagram is shown in Fig. 12 with a block diagram of the spike detection unit pictured in Fig. 13. The comparator array has as its inputs the time-division-mutiplexed data lines from the sixteen 2D PIA-2 probes that form the 3D array. The address counter provides timing information to the comparator array so that the comparators can be latched without disturbing the functionality of the circuit. The clocks of the 2D electrodes are also controlled by the address counter so that the input data can be synchnonized with the spike detection ASIC. When a spike is detected by the comparator array, the 7bit address of the spike along with a polarity bit and a control bit are stored in the register file. An input pointer keeps track of which register the data should be placed in. The register file allows for up to eight site address to be stored at a time to accomidate the case where several spikes are detected in succession. After being placed in the register file, the address and polarity information is transmitted serially off of the probes over one data lead. An output pointer controls which register is to be serially shifted off of the platform with the oldest spike data being shifted off first.

In order to provide a more through understanding of the spike detection unit operation, waveforms generated by the behavioral verilog are presented below. A simple case, where only one spike is present on one of the channels over the duration of the simulation is shown in Figs. 14 and 15. The operation is as follows. Initially, the ***reset*** data line goes low and the circuit begins operation. The data lines ***Comp_Inputs_N*** and ***Comp_Inputs_P***, represent the outputs from the comparators, which have yet to be designed. In Fig. 13, the input data is set up such that there is a positive spike present on channel one. When the address counter arrives at channel one, a spike is detected, ‘000000111’ is put in the first output register, and the input register pointer is incremented so that the address of the next spike detected will be placed in register two. In register one, bit zero is a control bit which allows for asychnonous transmission of the data off of the probes. Bit one is the polarity of the spike and bits two through eight are the site address. Since no other data is currently being shifted off of the probes, the data for the spike on channel one is serially transmitted off of the probe at the start of the clock cycle. The data is shifted out in reverse order, such that the control bit is shifted out first. When the control bit is high, it indicates to the external electronics that the next eight bits will be neural spike data. When all of the data in register one has been transmitted off of the platform, the contents of register one are reset to zeros and the output register pointer in incremented such that the next data to be serially shifted off of the probe will come from register two. The next time the spike detection unit scans address one there is no spike present on that channel. This situation is illustrated in Fig. 15. The spike detection unit works such that the address where a spike occurs will be transmitted off of the platform when the spike first occurs and when it ends. For all other times in between it is assumed that the spike is still present on the channel. This reduces the amount of data to be shifted off of the probes and prevents the transmission of bad data off of the probes in the case where a comparator or amplifier is railed to the positive or negative supply. Since in Fig. 15 there is no spike on channel one, and the time before when it was scanned there was a positive spike on channel one, the address for channel one and a negative spike polarity are placed in register two for transmission off of the platform.

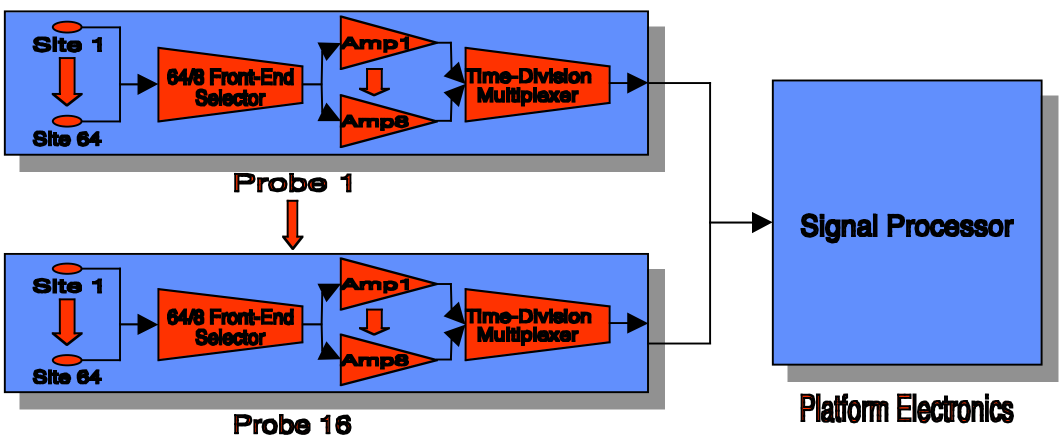

Fig. 12: System-level block diagram of a chronic recording system with on-platform spike recognition.

A more complicated scenerio where spikes are initially present on channels one and two is pictured in Figs. 16-18. When the spike detection unit scans addresses one and two, spikes are detected and the data from site one is placed into register zero, while the data from site address two is placed into register one. This data is then serially shifted off of the platform with the data from site one being transmitted first. Figure 17 shows the response of the spike detection unit the next time addresses one and two are scanned. Here, the spike on site two has ended, while the spike on site one is still present. Thus, the address of site two along with the polarity and control bits are placed in register two for transmission off of the platform. No data from site one will be transmitted because the neural spike is still present on that channel. The third time these two sites are scanned there are no spikes present on either channel as shown in Fig. 18. Since channel two had no spike present on it the last time it was scanned, no action is necessary in regard to this channel. The spike on channel one, however, has ended since the last time channel one was scanned. The response of the spike detection unit to this is to place the address, polarity, and control data from channel one into register three and then to serially shift this data to the outside world.

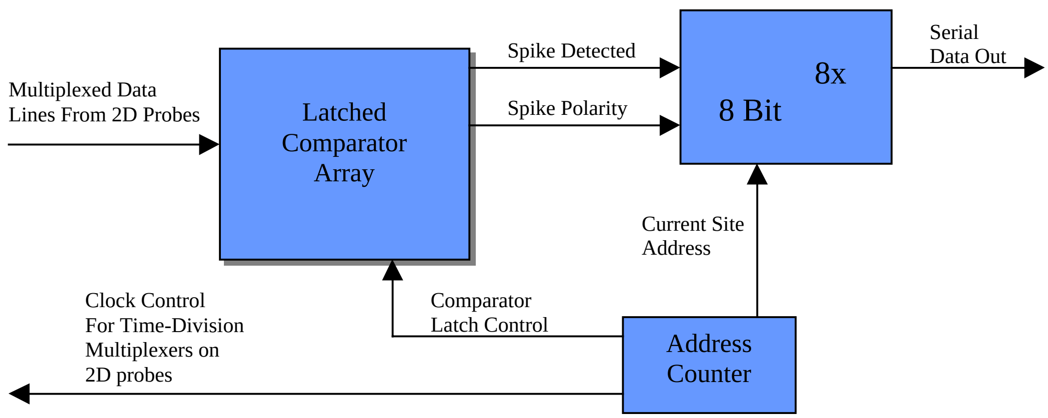

Fig. 13: Spike-detection unit block diagram

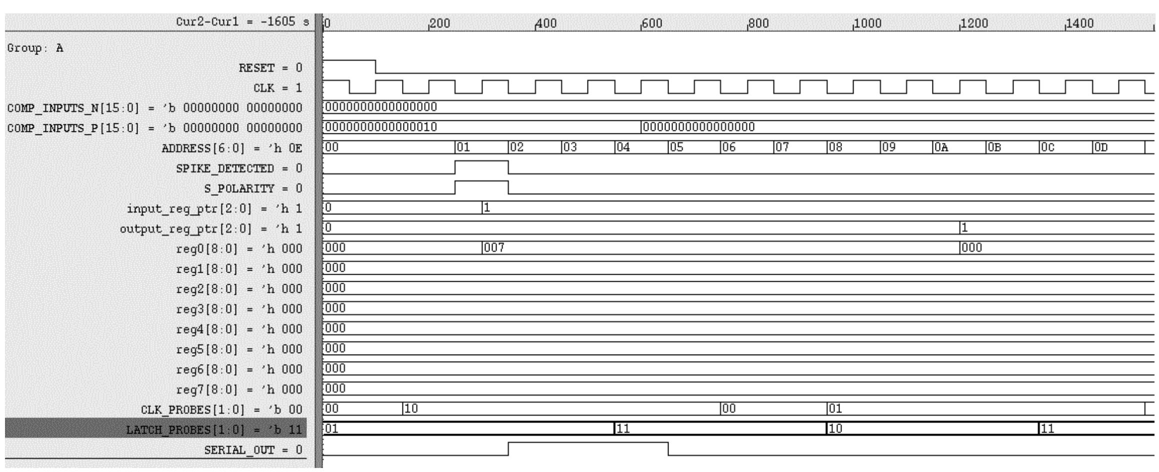

Fig. 14: Waveform data for a spike detected on channel one

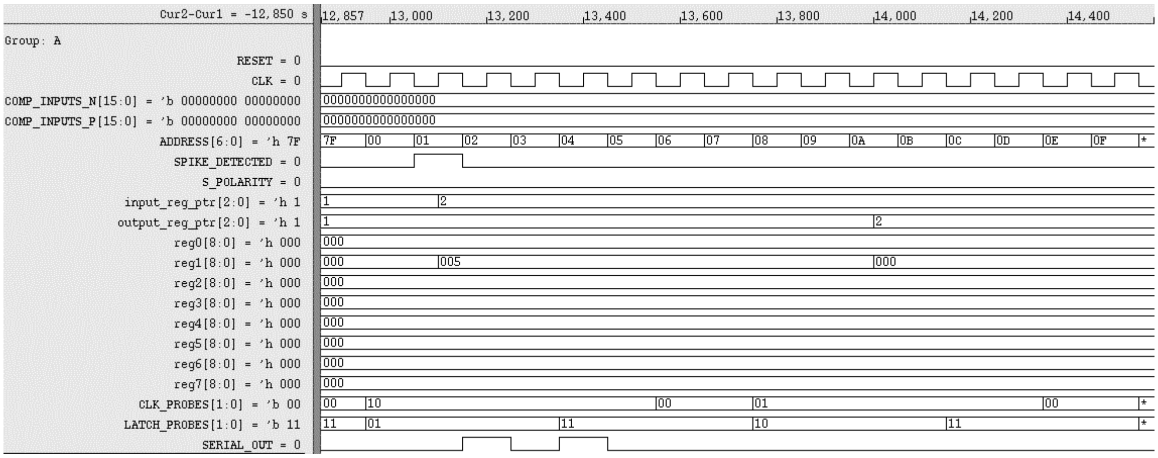

Fig. 15: Waveform data for a spike ending on channel one

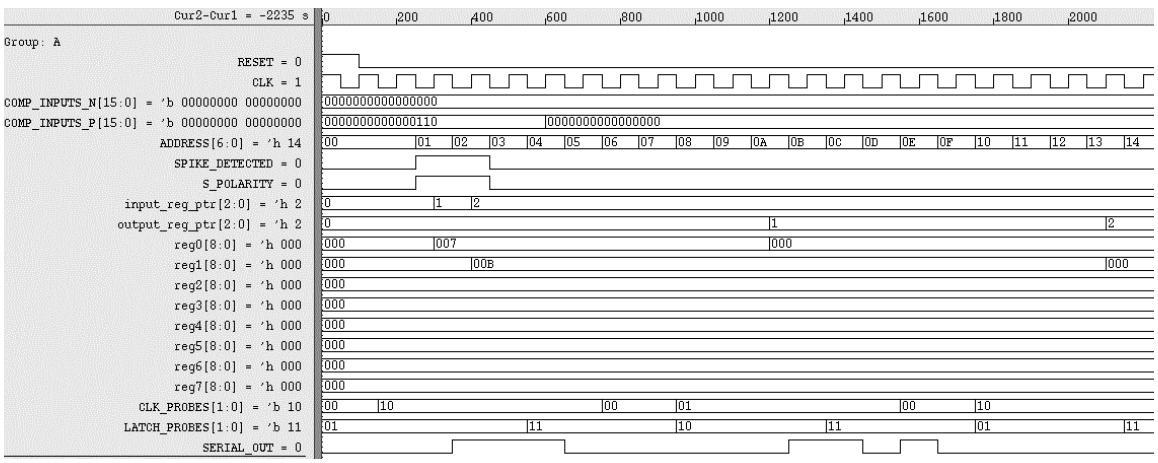

Fig. 16: Waveform data for spikes on channels one and two

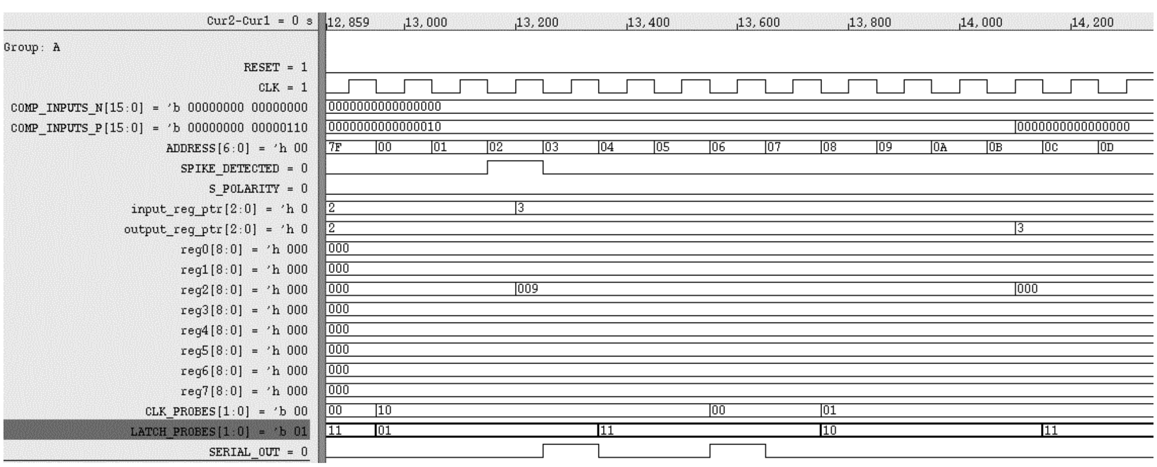

Fig. 17: Waveform data for a continued spike on channel one and a spike that has ended on channel two

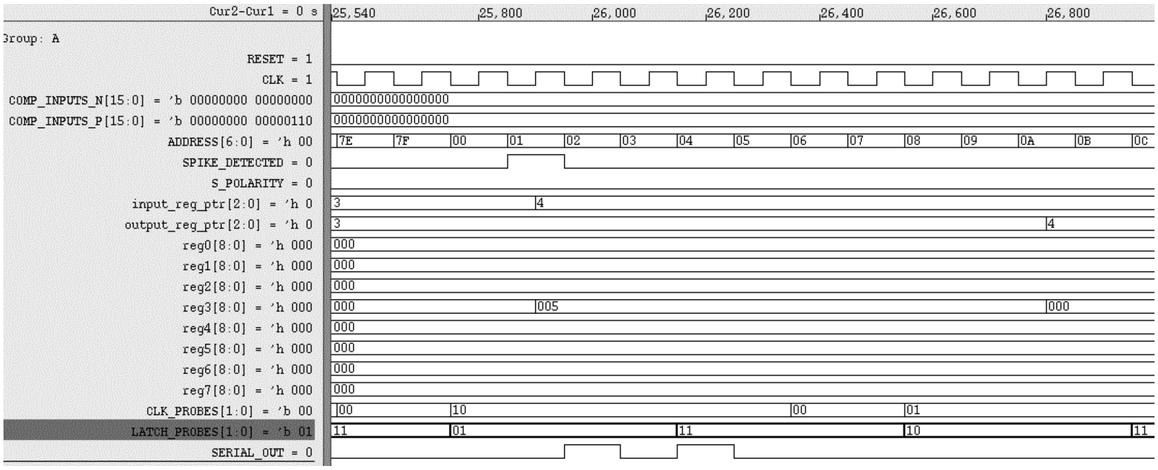

Fig. 18: Waveform data for a spike ending on channel one

5. Design of a Wireless Telemetry Platform for Multichannel Microprobes

As noted in earlier quarterly reports, we are developing a wireless interface for these probes to be mounted on the plaform. Testing of the latest telemetry chip was completed in this quarter. The chip design submitted to MOSIS for fabrication in late May and is shown in Fig. 19.

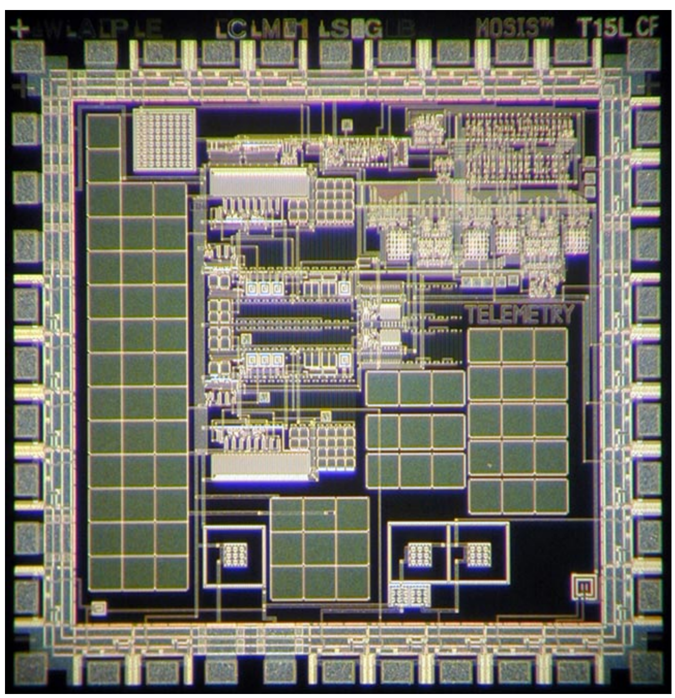

Fig. 19: The latest telemetry chip for chronic recording.

Some of the test results for this chip are listed and are analyzed below.

5.1 Front-End Circuitry

Compared with last fabricated chip, this design improves the front-end circuitry to achieve a lower power supply voltage and lower power consumption. The detailed testing of circuit blocks such as bandgap voltage reference (which have been verified on earlier chips) has so-far been bypassed in favor of the new more major circuit blocks. All four chips returned work well. There are three regulators on each chip. Two of the regulator output voltages are 3.3V; one is used for the digital circuit blocks and the other is used in the analog blocks. The other regulator output voltage is 1.65V, used in sigma-delta modulator as a mid-line. The line regulation of two regulators is shown in Fig. 20.

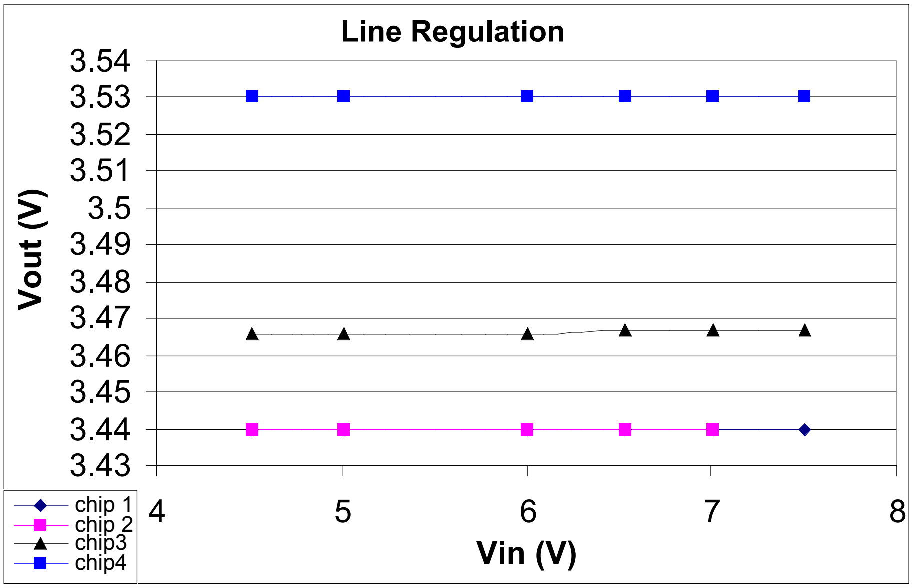

Fig. 20(a): Line regulation for the regulator with an output voltage of about 3.3V

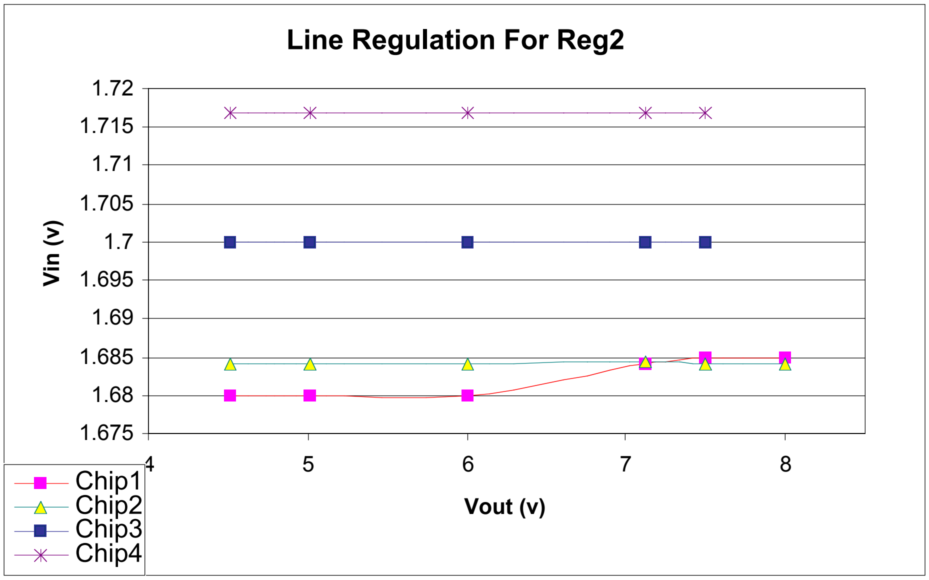

Fig. 20(b): Line regulation for the regulator with an output voltage of about 1.65V

From Fig. 20, the line regulation is better than 1mV/V. The ripple rejection ratio is shown in Fig. 21. It is observed that the ripple rejection ratio is better than 30dB. This performance may be worse than on the last chip, but the power consumption is much lower here than in that case. The total quiescent current of three regulators is about 30µA, but it is about 80µA on that chip.

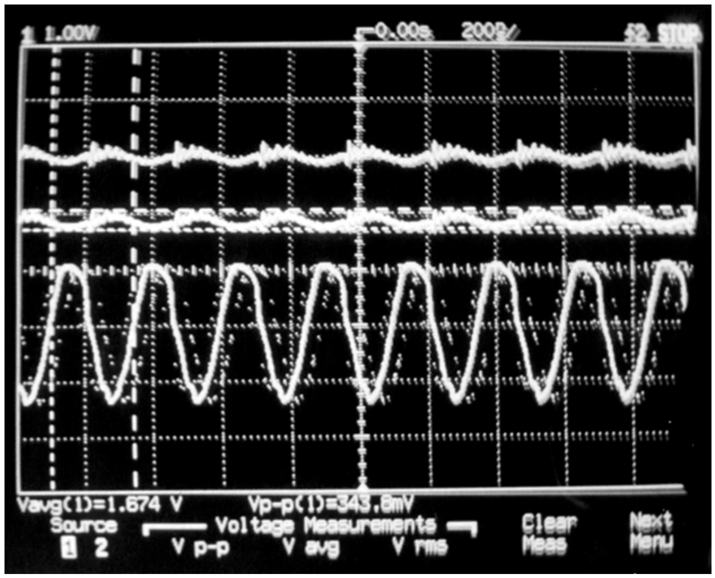

Fig. 21: The ripple rejection of the regulators.

The POR (power-on-reset) signal can be observed every time the chip is powered on, as shown in Fig. 22. The amplitude of the POR pulse is 3.4V and thepulse width is 75µs.

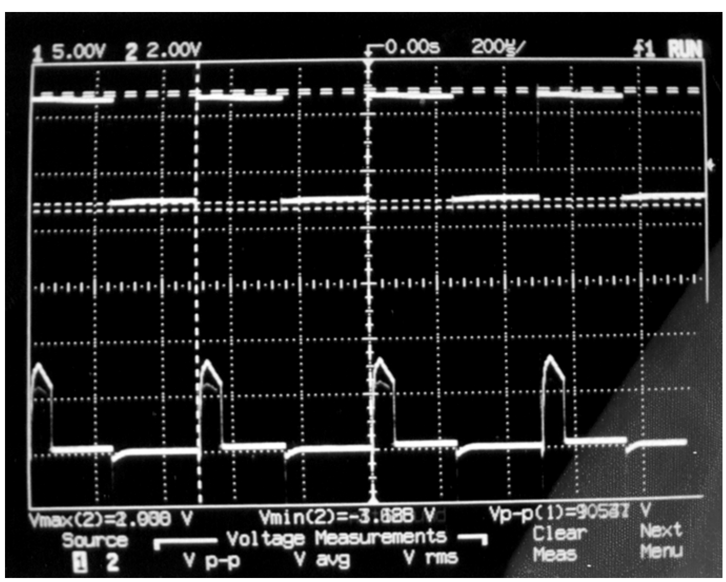

Fig. 22: The POR signal

5.2 Sigma-Delta Modulators

A 4MHz clock is recovered from the RF carrier. It is then divided by a D-flip-flop to 2MHz and used in the sigma-delta modulator (SDM). The recovered 2MHz clock is shown in Fig. 23. When shortening the two input terminals of the sigma-delta modulator, the direct output of the sigma-delta modulator is also shown in Fig. 23. The duty cycle of SDM output is 50%, which means that the average of the SDM is a constant of zero, which verifies the functionality of the SDM to some extent.

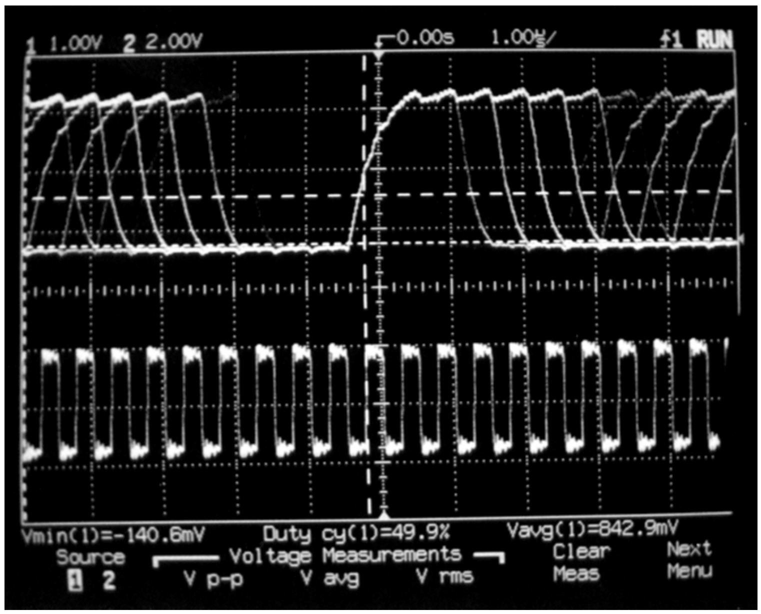

Fig. 23: The output of the sigma-delta modulator

The full functionality of the SDM can be verified by feeding the SDM with a pure sinusoid, recording the output, and then reconstructing the signal from the output with software or hardware. Recording the SDM output can be done using a DAQ (data acquisition) board and Labview. One of the recording processes is shown in Fig. 24. The maximum sample rate of the DAQ board is 500kHz, and it was used as the sample rate in recording the SDM output. This maximum sample rate limits the signals we can record because a certain oversampling ratio is required. Therefore, we chose to record the DSM output when the input was a pure sinusoid with a medium frequency of 1kHz. The results are found to be very good as seen in Fig. 25. A comparison between an ideal second-order SDM and the fabricated one can be found in Figs. 25 and 26. Figure 25 shows the reconstructed waveforms. Because the points in one cycle are not enough, it may not be sufficient only to compare these two waveforms. However, the spectra of the reconstructed signals in Fig. 26 are very helpful in showing that the eminent peaks appear at 1kHz, for both the ideal and the fabricated SDMs. This verifies the functionality of the fabricated SDM. We notice that the spectrum of the reconstructed waveform from the fabricated SDM has a very high DC component because the signal used to test the chip had a DC component of 1.8V.

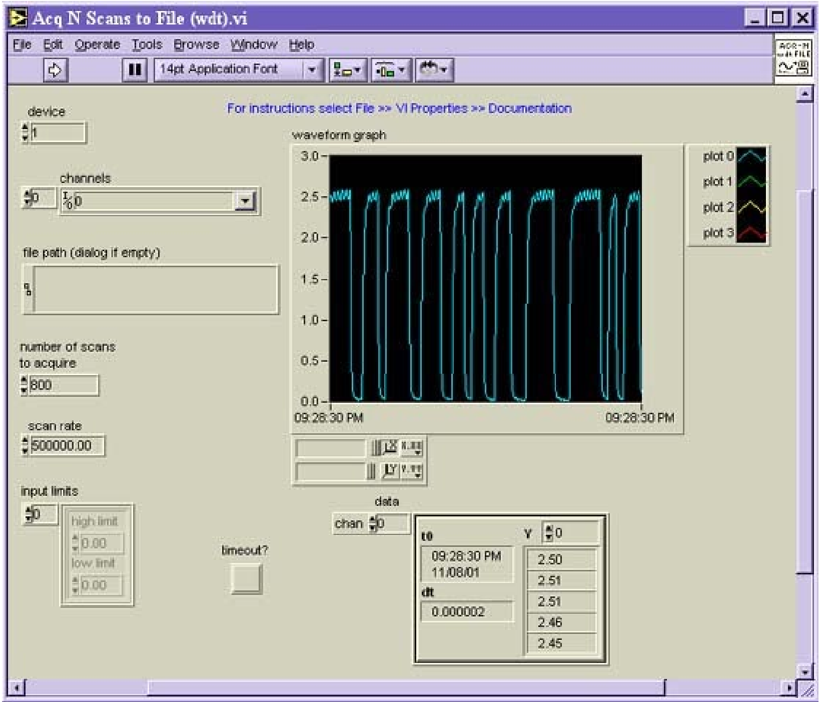

Fig. 24: The recorded SDM output waveform by Labview (sample rate = 500kHz)

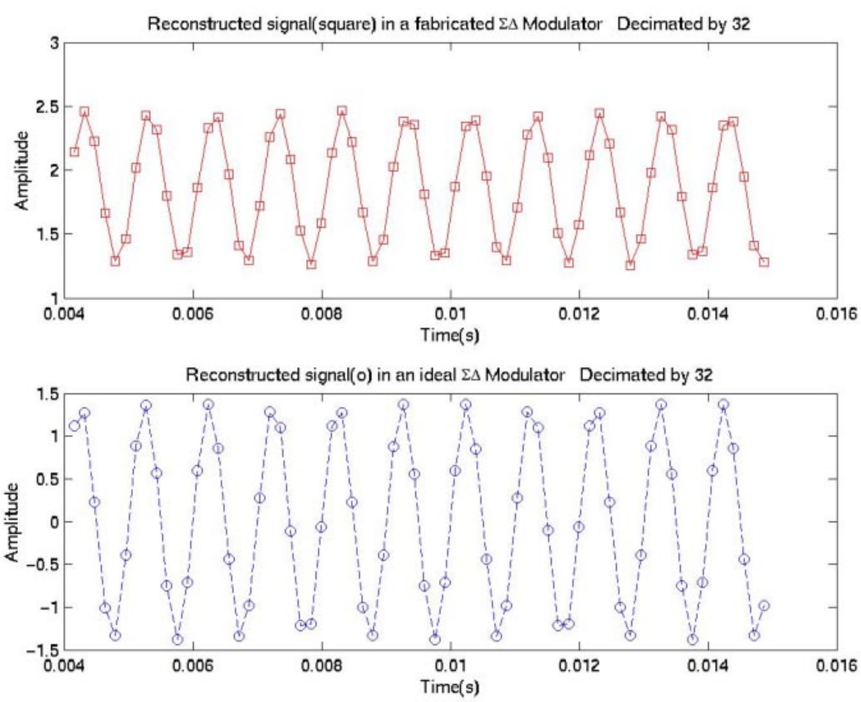

Fig. 25: The reconstructed sinusoids for an ideal and a fabricated SDM

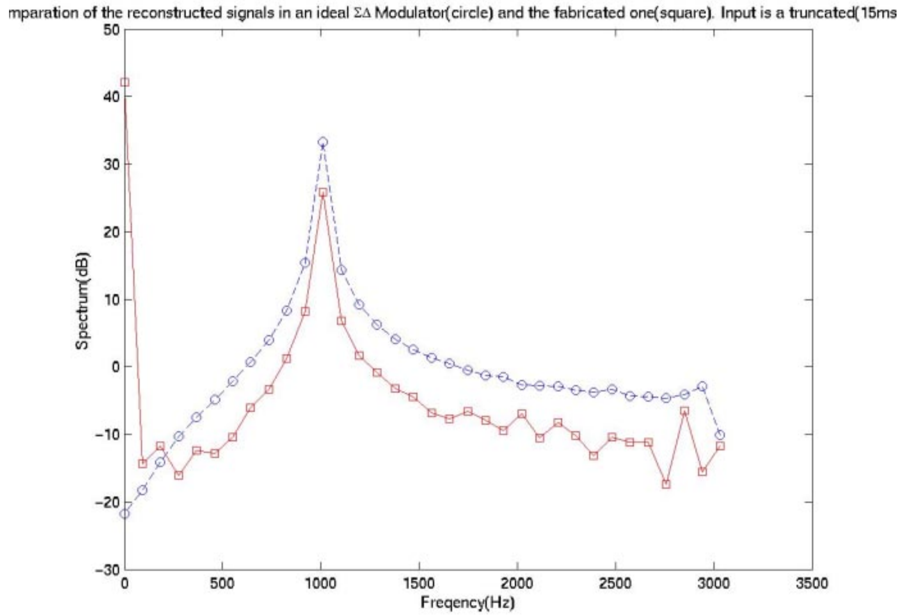

Fig. 26: The spectra of the reconstructed sinusoids for the ideal and fabricated SDMs

5.3 The On-Chip Transmitter

The on-chip transmitter will be used to transmit the digitized results to the external receiver, and the principle of data transmission used here is OOK (on-off-keying). The circuit diagram was shown in last quarterly report and is essentially a ring oscillator driving a transmit coil. The waveform of the oscillation is measured shown in Fig. 27.

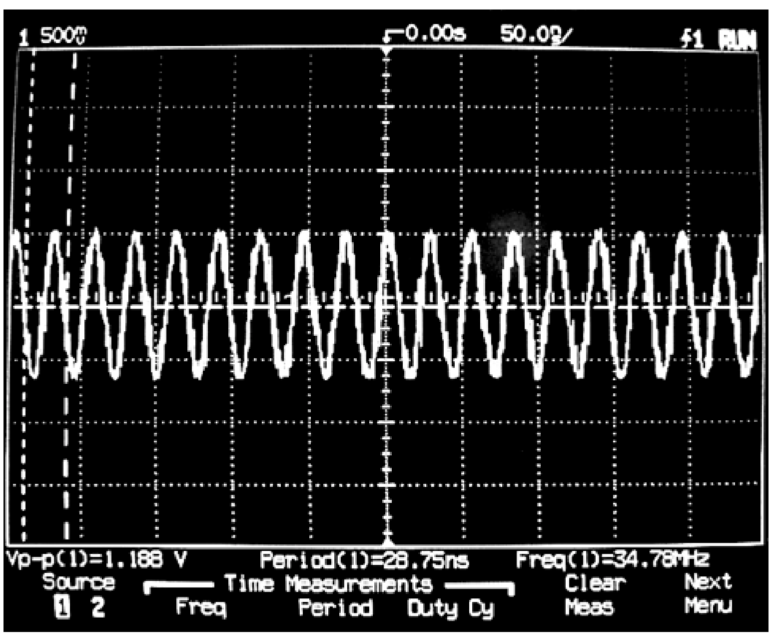

Fig. 27: The on-chip oscillator output

The frequency of the oscillator is 34.8MHz. This result is much lower than our expectation. The reason for the low output frequency is that the output of the ring oscillator is directly connected to the pad for testing, which introduces relatively large parasitic capacitance into the ring oscillator and slows down the oscillation. We expect output frequencies of more than 100MHz will be obtained when no testing pad is used.

A summary of the circuit performance from this chip is shown in Table 1 below.

Circuit Blocks	***Power Consumption***	***Specification***
Three Regulators; Two with Vout=3.3V; One for digital blocks and one for analog blocks One with V=1.6V	167µW when the received RF has a p-p amplitude of 12V.	VDD1 = 3.44V VDD2 = 1.68V Line Regulation < 1mV/V Ripple Rejection Ratio > 30dB
Clock Recovery	50µW	Clock Frequency = 4MHz
POR	No power consumption after POR signal	POR Pulse width = 75µs
Sigma Delta Modulator	~ 230µW when working at 2MHz	Output bit rate 2MHz
Ring Oscillator	~200µW	34.8MHz
Total Circuitry	<900µW 160µA When external power supply is 5.0V	Area = 2.2x2.2 mm^2

Table 1: Summary of the measured performance of the fabricated telemetry chip

6. Conclusions

During the past quarter, we have proceeded to develop chronically-implantable multi-channel recording probes in a number of areas. Work has begun to decrease the shank widths on the probes to much smaller dimensions (<20µm) while accommodating recording sites in the 600µm to 800µm range. A process based on polysilicon refill of a deep trench along the sides of the probe is being explored.

We have just completed the lastest fabrication run of active probes. Circuit thresholds are on-target at +0.9V (nMOS) and –0.6V (pMOS). Probes containing 3µm-thick gold shields over the circuitry have also been fabricated in preparation for chronic implant studies of the active structures. Some of the active probes contain silicon ribbon cables to allow them to be implanted without a surface platform, and for the first time such ribbons have been released entirely from the front using trenches formed along the sides of the cable. Also during the past quarter, the inputs from active probes have been stabilized with respect to their dc potentials using subthreshold shunt MOS transistor clamps. Without the clamps, the dc input level wanders ±1mV over short times and ±10-

40mV over longer periods. With the clamps activated, the dc input level is stable to better than ±50μV, while the lower ac cutoff frequency varies from about 30Hz to less than 3Hz, depending on bias. The use of a variable bias on the clamping devices would provide a means for varying the bandpass of the recording system depending on application.

With the dc baseline of the probes stable over time, it should be possible to implement systems in which the recorded spikes are thresholded and either the simple occurrence of a spike (together with its site address) are passed off the probe or a digital word giving its amplitude and address are transmitted. This can be expected to save large amounts of bandwidth compared to full analog signal transmission and sets the stage for prosthetic systems in which signal interpretation is done in-vivo. A platform chip performing spike thresholding has been designed and fully simulated to demonstrate its functional operation. The chip will be fabricated through MOSIS and used on platform-mounted 3D arrays of PIA-3 64-site 8-channel probes.

Finally, we have obtained the latest version of the telemetry interface back from fabrication and have begun testing it. The chip features a new front-end interface block together with a sigma-delta demodulator and a transmitter. The on-chip regulators consume 167μW for in input voltage of 12V and offer regulation of 1mV/V with a ripple rejection ratio of >30dB. The sigma-delta modulator consumes 230μW and produces an output bit rate at 2MHz. The total telemetry interface chip dissipates less than 900μW from 5V and measures 2.2mm x 2.2mm. We are moving to complete testing on both the PIA-2 chips and the telemetry interface and then hope to test the probes in a wireless configuration in-vivo.